arXiv:1405.6460v1 [stat.CO] 26 May 2014

Bayesian likelihood-free localisation of a biochemical source using multiple dispersion models

Branko Ristic*, Ajith Gunatilaka, Ralph Gailis, Alex Skvortsov
Land Division
Defence Science and Technology Organisation
Australia

May 27, 2014

Abstract

Localisation of a source of a toxic release of biochemical aerosols in the atmosphere is a problem of great importance for public safety. Two main practical difficulties are encountered in this problem: the lack of knowledge of the likelihood function of measurements collected by biochemical sensors, and the plethora of candidate dispersion models, developed under various assumptions (e.g. meteorological conditions, terrain). Aiming to overcome these two difficulties, the paper proposes a likelihood-free approximate Bayesian computation method, which simultaneously uses a set of candidate dispersion models, to localise the source. This estimation framework is implemented via the Monte Carlo method and tested using two experimental datasets.

1 Introduction

The threat of accidental or deliberate release of toxic (bio or chemical) aerosols into the atmosphere, has been well documented [1]. Wind, as the

*Corresponding author: B. Ristic, DSTO, LD, Bld 94, Melbourne, Vic 3207, Australia; email: branko.ristic@dsto.defence.gov.au; tel: +61 3 9626 8370; fax: +61 3 9626 8473

dominant transport mechanism in the atmosphere, can generate strong turbulent motion, causing the released aerosol to disperse as a plume whose spread increases with the downwind distance [2]. For the sake of public safety, it is of utmost importance to rapidly detect and localise a source of toxic release so that the mitigation actions can be carried out promptly. Ideally, the environment should be monitored continuously by a network of spatially distributed sensors to measure the concentration of toxic aerosols at various locations. This paper is devoted to the problem of source localisation using the concentration measurements collected by such a sensor network.

Two major difficulties are encountered in the described context of source localisation. The first is the choice of the most suitable dispersion model; the second is the lack of a precise and accurate probabilistic description of concentration measurements. Both difficulties are elaborated below.

A dispersion model describes, via mathematical equations, the physical processes that govern the atmospheric dispersion of the released material within the plume. The primary purpose of a dispersion model is to calculate the mean concentration of emitted material at given sensor locations. A plethora of dispersion models are in use today [3] to account for specific weather conditions, terrain, source hight, etc. The problem of selecting the most suitable dispersion model using statistical signal processing techniques [4] is largely neglected by the atmospheric research community. The only somewhat relevant reference is [5], which applies a single dispersion (Lagrange stochastic) model under multiple source assumptions in order to estimate the number of active sources of a toxic release.

Many references have been published on the topic of biochemical source localisation, under an adopted suitable dispersion model. The standard solutions are based on optimisation techniques, such as the nonlinear least squares [6], that are known to fail due to local minima or poor convergence. A version of the least-squares is used in [7] with minimisation carried out via simulated annealing. The alternatives are Bayesian techniques. Keats et al. [8] solved the chemical source localisation problem for the case of a transient release in the Bayesian framework using Markov chain Monte Carlo

(MCMC) and assuming Gaussian likelihood function of measurements. A similar approach was adopted in [9]. Bayesian framework with Gaussian likelihood of measurements was also applied in [10] to localise a biochemical source in urban environments. A version of the MCMC was proposed in [11], but assuming the log-normal likelihood function of measurements.

The Bayesian approaches, such as [8, 9, 10, 11, 5], are preferred than the optimisation techniques, because they result in the posterior density function of the source location, thereby providing an uncertainty measure to any point estimate derived from it. However, the Bayesian approaches referenced above, require a precise specification of the probabilistic model of the likelihood function (e.g. Gaussian, log-normal, with their parameters). Note that errors in the measurements are not only due to sensor noise but also due to the modeling inaccuracies, both of which in practice are very difficult to specify precisely. A large part of the difficulty arises from the fact that most approaches to source estimation have involved the use of ensemble mean concentration models of atmospheric tracer dispersion, but in reality, sensors are exposed to widely fluctuating concentration fields in both space and time, that are also stochastically non-stationary. This brings us to the second difficulty mentioned above: the lack of precise and accurate probabilistic descriptions of concentration measurements.

In this paper we develop a framework which overcomes both of the aforementioned difficulties: a likelihood-free approximate Bayesian computation (ABC) method for the localisation of a biochemical source in the combination with the selection of the most suitable model from a set of candidate dispersion models. This estimation framework is implemented via the Monte Carlo method and tested using two experimental datasets.

The paper is organised as follows. Section 2 describes the problem in a formal manner and also introduces a solution framework in the form of the multiple-model ABC rejection sampler. Section 3 presents the adopted dispersion models for source localisation. Section 4 describes the proposed adaptive iterative multi-model ABC sampler. Section 5 presents the numerical results, obtained using the experimental datasets. Finally, the main findings of this study are summarised in Section 6.

2 Background

2.1 Problem formulation

Let us assume that the system (the atmospheric tracer dispersion) obeys one of a finite number of models. A discrete random variable m denotes the model, whose domain $\mathcal{M} \subset \mathbb{N}$ has cardinality $M = |\mathcal{M}| \geq 1$. Each model $m \in \mathcal{M}$ is parameterised by a vector $\boldsymbol{\theta}_m \in \boldsymbol{\Theta}_m$, where $\boldsymbol{\Theta}_m \subseteq \mathbb{R}^m$ is the corresponding parameter space. The models, in general, are not nested [4], although the intersection of parameter spaces $\boldsymbol{\Theta} = \boldsymbol{\Theta}_1 \cap \cdots \cap \boldsymbol{\Theta}_M$ is non-empty; $\boldsymbol{\Theta}$ represents the space of the core parameters which includes the coordinates of source location.

Let a concentration measurement from sensor $s = 1, \ldots, S$ (sensor locations are known) be denoted ζ_s. All sensor measurements are stacked into a vector $\mathbf{z} = [\zeta_1, \ldots, \zeta_S]^\intercal$.

The problem is cast in the Bayesian framework. Let $\pi_{\boldsymbol{\theta}_m}(\boldsymbol{\theta}_m|m)$ denote the prior distribution over the parameter space $\boldsymbol{\Theta}_m$, with $m = 1, \ldots, M$. The posterior distribution of the parameter vector $\boldsymbol{\theta}_m$ follows from the Bayes theorem:

$$p(\boldsymbol{\theta}_m|m, \mathbf{z}) = \frac{\ell(\mathbf{z}|\boldsymbol{\theta}_m, m)\pi_{\boldsymbol{\theta}_m}(\boldsymbol{\theta}_m|m)}{p(\mathbf{z}|m)} \tag{1}$$

where $\ell(\mathbf{z}|\boldsymbol{\theta}_m, m)$ is the likelihood function and

$$p(\mathbf{z}|m) = \int_{\boldsymbol{\Theta}_m} p(\mathbf{z}|\boldsymbol{\theta}_m, m)\pi_{\boldsymbol{\theta}_m}(\boldsymbol{\theta}_m|m)d\boldsymbol{\theta}_m \tag{2}$$

is the marginal probability of measurement $\mathbf{z}$ given model m.

Let $\pi_m(m)$ denote the prior distribution over the candidate models $\mathcal{M}$. Then the posterior distribution over the models is also obtained using the Bayes theorem as:

$$p(m|\mathbf{z}) = \frac{p(\mathbf{z}|m)\pi_m(m)}{\sum_{m\in\mathcal{M}} p(\mathbf{z}|m)\pi_m(m)}. \tag{3}$$

Our ultimate goal is the posterior of the core parameter vector $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, which can be obtained via model averaging [12],

$$p(\boldsymbol{\theta}|\mathbf{z}) = \sum_{m\in\mathcal{M}} p(\boldsymbol{\theta}_m|m, \mathbf{z})p(m|\mathbf{z}), \tag{4}$$

where $p(\boldsymbol{\theta}|m, \mathbf{z})$ is the posterior (1) over the subspace $\boldsymbol{\Theta} \subseteq \boldsymbol{\Theta}_m$.

2.2 Approximate Bayesian computation

ABC constitutes a class of Bayesian-type algorithms developed for the estimation of a parameter vector in situations where the likelihood function is intractable or unknown [13]. For the case we consider, with multiple candidate models, the expressions for likelihoods $\ell(\mathbf{z}|\boldsymbol{\theta}_m, m)$, which feature in (1), are unknown. ABC methods replace the unknown likelihood $\ell(\mathbf{z}|\boldsymbol{\theta}_m, m)$ with the comparison between the observed measurement and the measurement synthesized using the model $m \in \mathcal{M}$. The simplest ABC algorithm is the ABC rejection sampler [14]. In the context of parameter estimation with multiple models, the ABC rejection sampler draws N samples from an approximation of the joint posterior $p(\boldsymbol{\theta}_m, m|\mathbf{z}) = p(\boldsymbol{\theta}_m|m, \mathbf{z})p(m|\mathbf{z})$. The pseudo-code of the ABC rejection sampler for multiple models is presented in Alg.1.

Algorithm 1 Multiple-model ABC rejection sampler

1: **Input**: $\mathbf{z}$; ϵ; N
2: Initialise: $\mathcal{X}_1 = \cdots = \mathcal{X}_M = \emptyset$
3: **repeat**
4: Draw m^* from $\pi_m(m)$
5: Draw $\boldsymbol{\theta}_m^*$ from $\pi_{\boldsymbol{\theta}_{m^*}}(\boldsymbol{\theta}_m)$
6: Simulate measurement $\mathbf{z}^*$ using model m^* and parameter $\boldsymbol{\theta}_m^*$
7: Compute distance $d^* = D(\mathbf{z}, \mathbf{z}^*)$
8: **if** $d^* \leq \epsilon$ **then**
9: $\mathcal{X}_{m^*} = \mathcal{X}_{m^*} \cup \{\boldsymbol{\theta}_m^*\}$
10: **end if**
11: **until** $\sum_{m=1}^{M} |\mathcal{X}_m| = N$
12: **Output**: $\mathcal{X}_1, \cdots, \mathcal{X}_M$

Distance $d = D(\mathbf{z}, \mathbf{z}^*)$ between the measurement vector and the synthesized data using model m^* is compared to the tolerance $\epsilon > 0$ in line 8 of Alg.1. The output of Alg.1 is a posterior $p(\boldsymbol{\theta}_m, m|d(\mathbf{z}, \mathbf{z}^*) \leq \epsilon)$ approximated

by M sets of random samples:

$$\mathcal{X}_m = \{\boldsymbol{\theta}_m^{(i)}\}_{1 \leq i \leq L_m}, \quad (m = 1, \ldots, M) \tag{5}$$

such that $\sum_{m=1}^{M} L_m = N$. Using (5) one can approximate the joint posterior $p(\boldsymbol{\theta}_m, m|\mathbf{z}) = p(\boldsymbol{\theta}_m|m, \mathbf{z})p(m|\mathbf{z})$ with:

$$p(m|\mathbf{z}) \approx \frac{L_m}{N} \tag{6}$$

$$p(\boldsymbol{\theta}_m|m, \mathbf{z}) \approx \frac{1}{L_m} \sum_{i=1}^{L_m} \delta(\boldsymbol{\theta}_m - \boldsymbol{\theta}_m^{(i)}) \tag{7}$$

where $\delta(x)$ is the Dirac delta function. The accuracy of approximation improves with larger N and smaller ϵ.

3 Dispersion models

This section describes $M = 3$ candidate dispersion models to be used in source localisation. The first two models are based on the Gaussian plume model [1, 2, 15] while the third model is referred to as the stretch exponential model [16, 17, 18].

Gaussian plume models adopt a Gaussian distribution of the plume in the vertical and horizontal directions under steady state conditions. These models are a solution of the equation of tracer transport with constant wind velocity (advection-diffusion equation). By convention, the wind velocity vector coincides with the x axis, while the spread of the plume in y and z directions is determined by the respective standard deviations σ_y and σ_z, commonly referred to as the Pasquill-Gifford sigmas [2, 15].

Consider a biochemical source of the release rate Q_0 located at coordinates (x_0, y_0, z_0). According to the Gaussian plume model, the mean concentration of the released material at the location of the sth sensor (x_s, y_s, z_s) is given by [2, 15]

$$C_s = \frac{Q_0}{2\pi\sigma_{y_s}\sigma_{z_s}U} \, e^{-\frac{(y_s-y_0)^2}{2\sigma_{y_s}^2}} \left[e^{-\frac{(z_s-z_0)^2}{2\sigma_{z_s}^2}} + e^{-\frac{(z_s+z_0)^2}{2\sigma_{z_s}^2}} \right] \tag{8}$$

if $x_s > x_0$ and zero otherwise. Notation U stands for the mean wind speed. Note that the Pasquill-Gifford sigmas, σ_{y_s} and σ_{z_s}, in (8) are assigned the sensor index s, because they are computed at coordinate x_s. The simplest model for σ_{y_s} and σ_{z_s} is based on the linear relationship with the downwind distance, i.e.

$$\sigma_{y_s} = \sigma_0 + \frac{\sigma_v}{U}(x_s - x_0), \tag{9}$$

$$\sigma_{z_s} = \sigma_0 + \frac{\sigma_w}{U}(x_s - x_0). \tag{10}$$

Explanation of the terms that feature in (9)-(10): σ_0 is the size of the source (in units of length); σ_v and σ_w are environmental parameters which account for the fluctuations in transverse and vertical velocities, respectively (in units of velocity). The first model we adopt for source localisation (referred to as $m = 1$) is the Gaussian plume model (8) with the spreads given by (9) and (10). The parameter vector for model $m = 1$ consist of 7 parameters, that is:

$$\boldsymbol{\theta}_1 = \begin{bmatrix} x_0 & y_0 & z_0 & \sigma_0 & B & \alpha & \beta \end{bmatrix}^\intercal, \tag{11}$$

where $B = Q_0/U$, $\alpha = \sigma_v/U$ and $\beta = \sigma_w/U$.

While the first dispersion model is very simple, it is unable to handle different canopy properties (average height, roughness, porosity). Model $m = 2$ is adopted to overcome this shortcoming. It is also based on (8) and (10), but the spread in y direction, σ_{y_s}, is a nonlinear function of the downwind distance from the source:

$$\sigma_{y_s} = \sigma_0 + \frac{\sigma_v}{U}\rho\left(\frac{x_s - x_0}{\rho}\right)^\gamma. \tag{12}$$

The exponent γ and the scale ρ (usually referred to as effective roughness) in (12) explicitly capture the canopy characteristics. The exponent γ may also change with meteorological conditions; its theoretical value for turbulent dispersion is $\gamma = 1/2$. The parameter vector for model $m = 2$ thus consists of 9 parameters:

$$\boldsymbol{\theta}_2 = \begin{bmatrix} x_0 & y_0 & z_0 & \sigma_0 & B & \alpha & \beta & \rho & \gamma \end{bmatrix}^\intercal. \tag{13}$$

The Gaussian plume dispersion models have two well-known deficiencies [2]. The assumption of constant wind speed does not even hold approximately for ground releases (i.e. $z_0 = 0$), and so vertical wind speed profiles are usually employed, such as the logarithmic wall-law from turbulent similarity theory, or the commonly used power-law profile [19, 20]. Use of these more physically realistic terms in the advection-diffusion equation imply fundamentally different solutions that are not of a Gaussian nature. More accurate analysis implies an "effective" plume convection velocity [18] that should be a function of downstream distance x_s, but this increases significantly the dimension of the parameter space. Also, under different meteorological conditions the functional form given by (8) can vary in the manner that expression for σ_{y_s} in (12) cannot fully capture.

The third dispersion model adopted for source localisation is the stretch exponential (SE) model. This model is included because it is capable, at least in theory, to overcome the aforementioned limitations of the Gaussian plume model. The SE model is a solution of the equation for tracer transport with a power-law wind velocity and turbulent diffusivity profiles, and is therefore more general that the Gaussian plume model. It enables an explicit, and rather simple, parameterisation of various meteorological condition and canopy characteristics. The pure SE model is only strictly applicable for ground-level sources. For $z > 0$, modified Bessel function solutions apply, which become difficult to handle for inverse source modelling applications.

For the current paper, we take a leading asymptotic term from these solutions, so that the mean concentration at the location of the sth sensor is given by [16, 17, 18]

$$C_s = \frac{B}{2\rho^2} \left(\frac{\rho}{x_s - x_0}\right)^{\tau} \left[\exp\left\{-\frac{z_s^r - z_0^r}{\sigma_z^r}\right\} + \exp\left\{-\frac{z_s^r + z_0^r}{\sigma_z^r}\right\}\right] e^{-\frac{(y_s - y_0)^2}{2\sigma_{y_s}^2}} \tag{14}$$

where

$$\sigma_{y_s} = \sigma_0 + \alpha\rho\left(\frac{x_s - x_0}{\rho}\right)^{1/2} \quad (15)$$

$$\sigma_{z_s} = \sigma_0 + \phi\rho r^{2/r}\left(\frac{x_s - x_0}{\rho}\right)^{1/r} \quad (16)$$

with $r = 1 + 2\mu$ and $\tau = 1/2 + (1 + \mu)/(1 + 2\mu)$ such that $0 \leq \mu \leq 1$. Parameters B, α and ρ have already been defined. Parameter ϕ is similar to parameter β in (11) and (13). Parameter r is introduced to capture the variability of meteorological conditions. It defines the functional form of the vertical concentration profile: for $r = 1$, this profile is exponential, while for $r = 2$ it is Gaussian. Parameter τ describes the mean concentration decay along the plume centreline and captures the variability of meteorological conditions (but is also affected by the type of canopy). According to the SE model, r and τ are related (both depend on μ). However, this relationship is valid only for an idealized flow with a power-law profile over the flat underlying surface (in this case μ is the exponent in the wind velocity profile). In order to make the model more flexible, we adopt a relationship $\tau = \nu + (1 + \mu)/(1 + 2\mu)$, where ν is a free parameter whose prior probability has the mean value equal to 1/2. Equation (14) is the first term of an expansion of the exact solution for the mean concentration at downwind distance, much greater than the plume spread.

In summary, the SE model consists of 10 parameters; its parameter vector is specified as follows:

$$\boldsymbol{\theta}_3 = \left[x_0 \quad y_0 \quad z_0 \quad \sigma_0 \quad B \quad \alpha \quad \phi \quad \rho \quad \mu \quad \nu\right]^{\mathsf{T}}. \quad (17)$$

As it was mentioned in Sec.2.1, the parameter spaces of the three models are not nested.

4 Adaptive iterative multiple-model ABC sampler

The ABC rejection sampler described in Alg.1 is very inefficient due to its low acceptance rate. Several improvements of the ABC rejection have

been proposed in the single model case, such as the ABC MCMC sampler [21] and a few versions of the ABC sequential Monte Carlo (SMC) sampler [22, 23, 24]. Following [23], we propose an iterative Monte Carlo multiple-model ABC sampler, whose basic steps are described by Alg. 2. The key feature of this algorithm is that it performs the ABC rejection scheme using a monotonically decreasing sequence of tolerance levels $\epsilon_1 > \epsilon_2 > \ldots \epsilon_{\mathrm{T}} \geq 0$, until the final tolerance is reached. The sequence of tolerances ϵ_t, $t = 1, 2 \ldots, T$ is computed by the algorithm, from the measurement vector $\mathbf{z}$, hence the number of iterations T is not know in advance. For a given model m, sampling is initially carried out from the prior $\pi_{\boldsymbol{\theta}_m}(\boldsymbol{\theta}_m)$, followed by sampling from a sequence of intermediate distributions $p(\boldsymbol{\theta}_m|m, d(\mathbf{z}, \mathbf{z}^*) \leq \epsilon_t)$, $t = 1, 2, \ldots,$ which gradually approach the target distribution $p(\boldsymbol{\theta}_m|m, d(\mathbf{z}, \mathbf{z}^*) \leq \epsilon_{\mathrm{T}})$. The theoretical justification of the proposed iterative scheme, presented in [23], is based on the sequential importance sampling (SIS) paradigm.

There are two differences between our proposed Alg. 2 and the algorithm reported in [23], both of which reflect the adaptive nature of the former: (i) our algorithm does not require the sequence of tolerances $\epsilon_1, \epsilon_2, \cdots, \epsilon_{\mathrm{T}}$ to be specified as an input (it works it out from the data $\mathbf{z}$); (ii) The proposal distribution for each sample and at each iteration is computed adaptively, in a manner similar to population Monte Carlo techniques [25].

Algorithm 2 : Adaptive iterative multiple-model ABC rejection sampler

1: **Input**: $\mathbf{z}$; Δ; N;
2: $[\{\mathcal{X}_m^0\}_{1\leq m\leq M}, \epsilon_1]$ = Init-Iter($\mathbf{z}, N$)
3: $t = 0$; $\epsilon_0 \leftarrow \infty$
4: **while** $(\epsilon_t - \epsilon_{t+1}) > \Delta$ **do**
5: $\quad t = t + 1$
6: $\quad [\{\mathcal{X}_m^t\}_{1\leq m\leq M}, \epsilon_{t+1}]$ = Repeated-Iter($\{\mathcal{X}_m^{t-1}\}_{1\leq m\leq M}, \epsilon_t, \mathbf{z}, N$)
7: **end while**
8: **Output**: $\{\mathcal{X}_m^t\}_{1\leq m\leq M}$

The initial iteration of Alg. 2 (line 2) creates M initial sample sets $\{\mathcal{X}_m^0\}_{1\leq m\leq M}$ from the priors π_m, $\pi_{\boldsymbol{\theta}_m}$, $m = 1, \ldots, M$, and computes the

first tolerance level ϵ_1. The repeated iteration, line 6, performs the rejection sampling at a given tolerance and also computes the tolerance level for the next iteration. The "while loop" (lines 4-7) is terminated when the difference between the two consecutive tolerances is below a certain threshold Δ.

Since the algorithm is based on SIS, both the initial and the subsequent samples $\{\mathcal{X}_m^t\}_{1\leq m\leq M},\ t=0,1,..$ are *weighted*. Hence (5) now takes the form:

$$\mathcal{X}_m^t = \left\{\left(w_m^{(i,t)}, \boldsymbol{\theta}_m^{(i,t)}\right)\right\}_{1\leq i\leq L_m^t}, \qquad (m=1,\ldots,M; t=0,1,\ldots) \tag{18}$$

where importance weights $w_m^{(i,t)}$ are normalised, that is $\sum_{i=1}^{L_m^t} w_m^{(i,t)} = 1$.

The steps of the initial iteration are given by Alg.3. The tolerance ϵ_1 is computed in line 10 as an order statistic φ of sample distances $d_1, \cdots, d_N$.

Algorithm 3 : $[\{\mathcal{X}_m^0\}_{1\leq m\leq M}, \epsilon_1] = \text{Init-Iter}(\mathbf{z}, N)$

1: **Input**: $\mathbf{z}$; N
2: $\mathcal{X}_1^0 = \cdots = \mathcal{X}_M^0 = \emptyset$
3: **for** $i = 1,\ldots,N$ **do**
4: Draw $m \sim \pi_m$
5: Draw $\boldsymbol{\theta}^* \sim \pi_{\boldsymbol{\theta}_m}$
6: Simulate measurement $\mathbf{z}^*$ using model m with parameter $\boldsymbol{\theta}^*$
7: Compute distance $d_i = D(\mathbf{z}, \mathbf{z}^*)$
8: $\mathcal{X}_m^0 = \mathcal{X}_m^0 \cup \{(w=1, \boldsymbol{\theta}^*)\}$
9: **end for**
10: $\epsilon_1 = \varphi(d_1, \ldots, d_N)$
11: **for** $j = 1,\ldots,M$ **do**
12: Normalise weights in $\mathcal{X}_j^0$
13: **end for**
14: **Output**: $\mathcal{X}_1^0, \cdots, \mathcal{X}_M^0$, ϵ_1

Finally, the pseudo-code of the repeated iteration is given in Alg.4. The while-loop between lines 4 to 17 carries out rejection sampling. The model m is drawn from the prior in line 5, followed by the selection of the sample index k in line 7. This selection is based on the weights in $\mathcal{X}_m^{t-1}$. A candidate

sample $\boldsymbol{\theta}^*$ is drawn from the proposal distribution $q_m(\theta_m|\boldsymbol{\theta}_m^{(k,t-1)})$ in line 8. The proposal is adopted as a normal distribution, whose mean is $\boldsymbol{\theta}_m^{(k,t-1)}$ and the covariance matrix is related to the sample covariance of the sample $\mathcal{X}_m^{t-1}$. The practical implementation of line 8 is as follows: $\boldsymbol{\theta}^* = \boldsymbol{\theta}_m^{(k,t-1)} + \lambda h \Sigma \varepsilon$, where $0 < \lambda < 1$ is a parameter of the algorithm, h is the optimal bandwidth of the multivariate Gaussian kernel [26, Eq.(12.2.7)], Σ is the square-root of the empirical covariance matrix of $\mathcal{X}_m^{t-1}$ and ε is a sample from the standard normal distribution. The computation of an unnormalised weight for an accepted sample in line 12, is carried out in the manner of [23]. The output tolerance ϵ_{t+1} is again computed as an order statistic φ of the sample $d_1, \ldots, d_N$.

Recall that out goal is to localise the source. Denote by vector $\ell = [x_0 \;\; y_0]^\intercal \in \mathbb{L}$ the coordinates of the source. The space of source coordinates is clearly a subspace of the core parameter space. From the output of Alg.2, expressed by (18) at the last iteration $t = T$, we can extract M sets of random samples over the subspace $\mathbb{L}$:

$$\mathcal{L}_m^t = \left\{ \left(w_m^{(i,t)}, \ell_m^{(i,t)} \right) \right\}_{1 \le i \le L_m^t}, \qquad (m = 1, \ldots, M; t = T). \tag{19}$$

According to (4), the posterior density of source location is then approximated by $\mathcal{L}_m^t$ as follows:

$$p(\ell|\mathbf{z}) \approx \sum_{m=1}^{M} \frac{L_m}{N} \sum_{i=1}^{L_m} w_m^{(i,t)} \, \delta(\ell - \ell_m^{(i,t)}) \tag{20}$$

5 Numerical analysis

5.1 Experimental datasets

Algorithm evaluation was carried out using two experimental datasets collected by COANDA Research & Development Corporation. The experiments were carried out using their large recirculating water channel, specially designed for dispersion modelling. The water channel is 10 m long, 1.5 m wide and 0.9 m deep. The floor of the water channel was covered with a metal mesh of height 4 mm to give surface roughness.

Algorithm 4 :$[\{\mathcal{X}_m^t\}_{1\leq m\leq M}, \epsilon_{t+1}] = \text{Repeated-Iter}(\{\mathcal{X}_m^{t-1}\}_{1\leq m\leq M}, \epsilon_t, \mathbf{z}, N)$

Input: $\{\mathcal{X}_m^{t-1}\}_{1\leq m\leq M}$, ϵ_t, $\mathbf{z}$, N
$\mathcal{X}_1^t = \cdots = \mathcal{X}_M^t = \emptyset$
$n = 0$
while $n < N$ **do**
    Draw $m \sim \pi_m$
    $L_m = |\mathcal{X}_m^{t-1}|$
    Select index $k \in \{1, \ldots, L_m\}$ with $\mathbb{P}(k = j) = w_m^{(j,t-1)}$
    Draw $\boldsymbol{\theta}^* \sim q_m(\cdot|\boldsymbol{\theta}_m^{(k,t-1)})$
    Simulate measurement $\mathbf{z}^*$ using model m with $\boldsymbol{\theta}^*$
    Compute distance $d^* = D(\mathbf{z}, \mathbf{z}^*)$
    if $d^* \leq \epsilon_t$ **then**
        Weight $\tilde{w}^* = \frac{\pi_{\boldsymbol{\theta}_m}(\boldsymbol{\theta}^*)}{\sum_{i=1}^{L_m} w_m^{i,t-1} q_m(\boldsymbol{\theta}^*|\boldsymbol{\theta}_m^{(i,t-1)})}$
        Accept: $\mathcal{X}_m = \mathcal{X}_m \cup \{(\tilde{w}^*, \boldsymbol{\theta}^*)\}$
        $n = n + 1$
        $d_n = d^*$
    end if
end while
$\epsilon_{t+1} = \varphi(d_1, \cdots, d_N)$
for $m = 1, \ldots, M$ **do**
    Normalise weights in $\mathcal{X}_m^t$
end for
Output: $\{\mathcal{X}_m^t\}_{1\leq m\leq M}$, ϵ_{t+1}

The source was releasing fluorescein dye, at a constant rate, from a narrow vertical tube, placed $z_0 = 4$ mm above the bottom of the channel at coordinates $x_0 = -373.5$ mm, $y_0 = 0$ mm. Concentration data were collected at several downstream positions using 1D laser induced fluorescence linescan system, at the rate of 300 lines per second, for a total sampling time of 1000 seconds [27].

Two experimental datasets are used for algorithm evaluation (both are available as supplementary material of this submission). Dataset 1 was collected in the absence of any obstacles (mimicking an open terrain scenario). Dataset 2 was collected in the presence of 10 mm high obstacles, placed on a regular grid, thus mimicking an urban scenario. Both datasets were extracted from the full recordings averaged over 100 seconds, and consist of $S = 48$ sensor measurements (four rows of 12 sensors) at downstream positions. The top-down view of the experimental setup for both cases is shown in Fig.1. The source location at coordinates $(-373.5, 0)$ is marked by a red asterisk. The position of sensors is indicated by blue circles whose radius is proportional (on the log-scale) to the corresponding concentration measurement. The height of all sensors in both setups was $z_s = 9.3$ mm, for $s = 1, \ldots, 48$.

5.2 Priors and parameters

We adopted the prior over $M = 3$ models to be uniform, that is $\pi_m(1) = \pi_m(2) = \pi_m(3) = 1/3$.

The priors for various parameters included in vectors $\boldsymbol{\theta}_1$, $\boldsymbol{\theta}_2$ and $\boldsymbol{\theta}_3$, specified by (11), (13) and (17), respectively, were adopted as follows (all units of length are *millimeters*):

$$\begin{aligned}
\pi(x_0) &= \mathcal{U}(-1000, 0), & \pi(y_0) &= \mathcal{U}(-500, 500),\\
\pi(z_0) &= \mathcal{G}(1.333, 3), & \pi(\sigma_0) &= \mathcal{G}(15.5, 0.03),\\
\pi(B) &= \mathcal{G}(2, 2.5), & \pi(\alpha) &= \mathcal{G}(3, 0.5),\\
\pi(\beta) &= \mathcal{G}(1.667, 0.15), & \pi(\rho) &= \mathcal{G}(6, 1),\\
\pi(\phi) &= \mathcal{G}(1.667, 0.15), & \pi(\gamma) &= \mathcal{B}(3, 3),\\
\pi(\mu) &= \mathcal{B}(1.5, 3), & \pi(\nu) &= \mathcal{B}(6, 6).
\end{aligned}$$

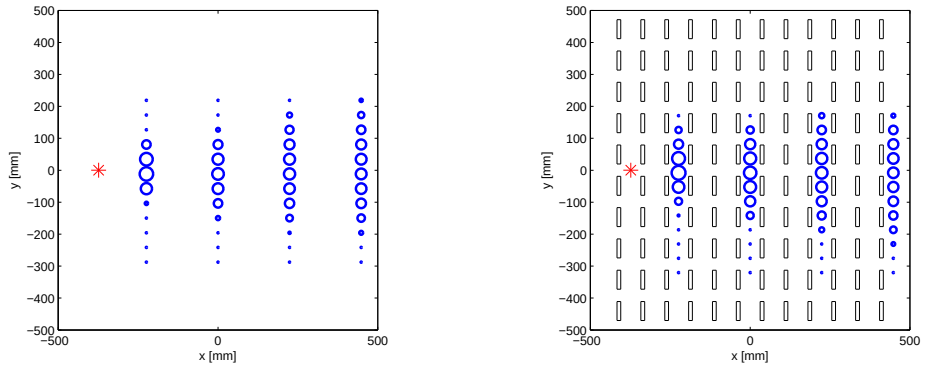

(a) (b)

Figure 1: Top-down view of the experimental setup: (a) dataset 1, (b) dataset 2. Location of the source is marked by a red asterisk at $(-373.5, 0)$. Sensor locations indicated by blue circles, whose radius is proportional (on the log-scale) to the corresponding concentration measurement. The rectangles in (b) indicate the contours of the obstacles

Here $\mathcal{U}(a, b)$ is the uniform distribution, with limits a and b, $\mathcal{G}(k, \eta)$ is the Gamma distribution with shape k and scale η, and $\mathcal{B}(p, q)$ is the Beta distribution with parameters p and q.

The proposed ABC sampler described by Alg.2, was executed using $N = 1000$ samples. The distance D between the actual measurement $\mathbf{z} = [\zeta_1, \ldots, \zeta_S]^\intercal$ and the synthesised "measurement", using model m, denoted $\mathbf{z}^* = [C_1(\boldsymbol{\theta}_m), \ldots, C_S(\boldsymbol{\theta}_m)]^\intercal$, was adopted as:

$$d(\mathbf{z}, \mathbf{z}^*) = \sum_{s=1}^{S} (\zeta_s - C_s(\boldsymbol{\theta}_m))^2. \tag{21}$$

Statistics φ, which computes the tolerance level for the next iteration, was adopted to be the 128th smallest value of samples $d_1, \ldots, d_{N=1000}$. The parameter λ, used in the proposal q_m, was set to $\lambda = 0.4$. The termination threshold was adopted as $\Delta = 2 \cdot 10^{-10}$.

5.3 Results

The results for dataset 1 are presented first. The tolerance levels computed by the proposed ABC sampler are shown in Fig.2.(a). It took $t = 10$ iterations to reach the final tolerance level of $\epsilon_{\mathrm{T}} = 4.04 \cdot 10^{-9}$. The average acceptance rate, over all iterations, was 2.8%. The model probabilities, $p(m|\mathbf{z})$ are shown at each iteration in Fig.2.(b). After the initial iteration ($t = 0$), the probabilities of all three models are approximately 1/3. Subsequently, while the tolerance levels were high, the three probabilities fluctuate until the iteration $t = 7$. After that, the probability of model $m = 1$ drops to zero, while the probability of models $m = 2$ grows to about 0.85. The probability of model $m = 3$ drops, but never goes to zero.

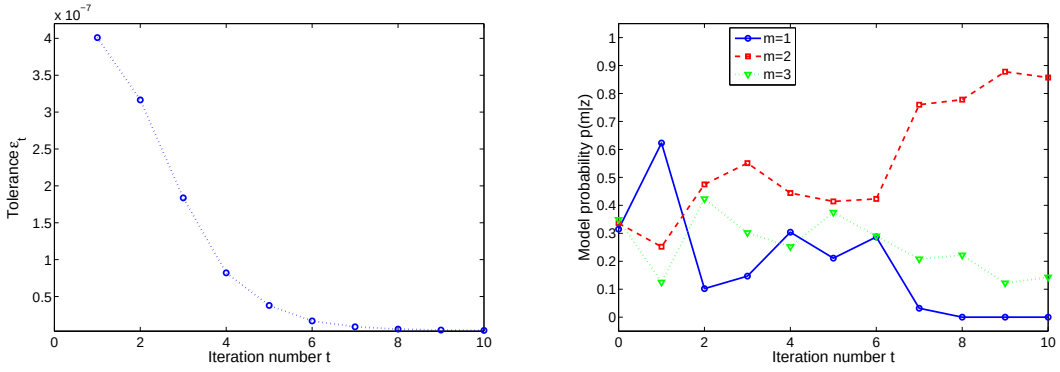

(a) (b)

Figure 2: Tolerance levels and model probabilities over iterations (dataset 1): (a) ϵ_t; (b) $p(m|\mathbf{z})$.

Fig.3 shows the scatter plot of random samples $\mathcal{L}_m^t$ which approximate the posterior density $p(\ell|\mathbf{z})$ of (20), at iterations (a) $t = 1$, (b) $t = 4$, (c) $t = 7$ and (d) $t = 10$. The true source location is marked by a red asterisk. Localisation of the source clearly improves with iterations, as the tolerance levels get smaller.

Finally, Figs.4.(a) and (b) display the final posterior distribution after the iteration $t = 10$, marginalised to x, and y axes, respectively. The density estimates from the random sample are obtained using the kernel density

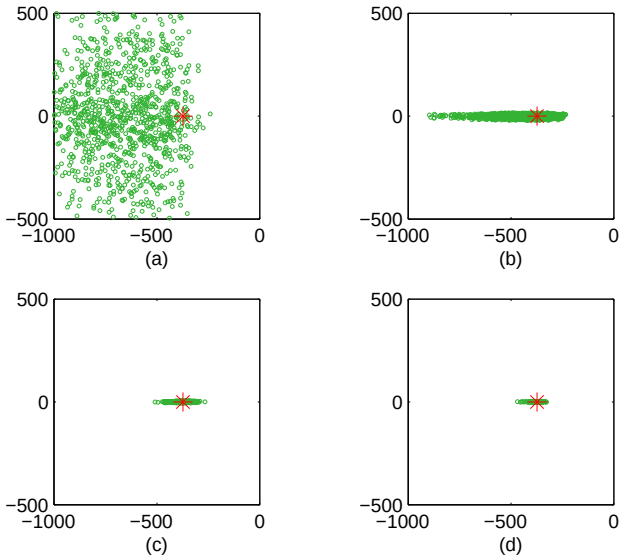

Figure 3: Scatter plots of random samples $\mathcal{L}_m^t$ in the (x, y) plane (dataset 1). The samples approximate the posterior density $p(\ell|\mathbf{z})$. Scatter plots shown after iteration: (a) $t = 1$, (b) $t = 4$, (c) $t = 7$, (d) $t = 10$. The true source location is marked with a red asterisk.

estimation (KDE) [28]. The true source coordinates are marked by the solid vertical lines.

The first observation is that the estimated marginal posteriors are remarkably accurate (considering that the prior was a uniform density with the span of 1000 mm) and include the true source location. The second observation is that the estimation is significantly better for the coordinate y_0, than for x_0, which is in accordance with the theoretical analysis carried out using the Carmer-Rao bound [29].

The dataset 2 was collected in a more challenging setup with obstacles. The results for dataset 2 are shown in Figs. 5, 6 and 7. Fig.5.(a) shows the computed tolerance levels over all iterations: note that it took $t = 12$ iterations to reach the final tolerance level of $\epsilon_{\mathrm{T}} = 2.75 \cdot 10^{-9}$. The average acceptance rate, for all iterations, was 2.9%. The model probabilities,

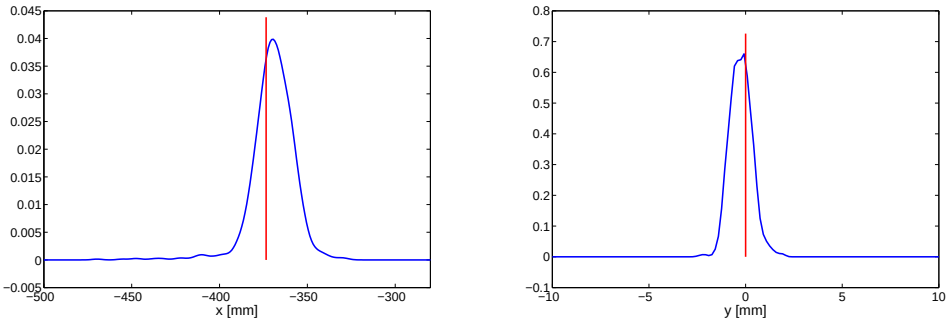

(a) (b)

Figure 4: Estimated posterior density of the source location (dataset 1): (a) x axis, (b) y axis. The true values are indicated by the vertical solid lines.

$p(m|\mathbf{z})$ are plotted in Fig.5.(b). We can observe that after iterations 7,8 and 9, $m = 1$ appears to be the preferred model. However, at iteration 10, its probability drops to zero. As in the case of dataset 1, the probabilities of models $m = 2$ and $m = 3$ remain non-zero at the lowest tolerance levels. It has already been noted that stretched exponential solutions may not be the most appropriate model for these datasets, even though the model has some physical motivation [27]. The analysis presented here reinforces these findings, and highlights the fact that more complex, higher-dimensional models are not necessarily better in describing inherently stochastic phenomena.

Fig.6 displays the scatter plot of random samples $\mathcal{L}_m^t$ at iterations: (a) $t = 3$, (b) $t = 6$, (c) $t = 9$ and (d) $t = 12$. The true source location is marked by a red asterisk. Once again, localisation of the source improves with iterations. However, a close inspection of the final localisation posterior $p(\ell|\mathbf{z})$, shown in Fig.6.(d), reveals that this posterior is bi-modal. This can be seen from Figs.7 which displays the final posterior distribution marginalised to x, and y axes, respectively. The two modes, clearly seen in Fig.7.(a), correspond to models $m = 2$ and $m = 3$, after iteration $t = 12$. They both appear to be biased: the stronger mode is due to model $m = 3$ (the probability of this model is higher at $t = 12$, see see Fig.5.(b)); it peaks approximately at $\hat{x}_0 \approx -350$.

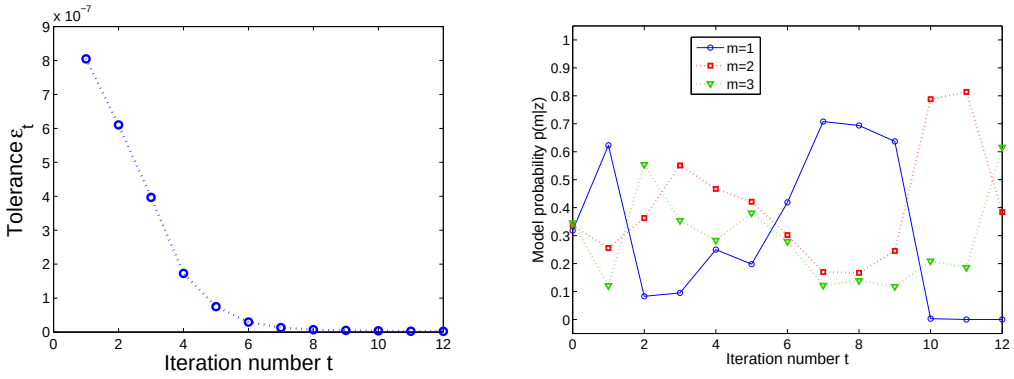

(a) (b)

Figure 5: Tolerance levels and model probabilities over iterations (dataset 2): (a) ϵ_t; (b) $p(m|\mathbf{z})$.

We can make the following observations with respect to the results obtained using dataset 2. First, none of the models seem to be correct since the estimation of the x_0 coordinate appears to be biased. Second, the use of multiple-models was beneficial, because the support of the posterior contains the true source location (albeit in the tail of one of the modes).

Finally, a remark on the proposed iterative multiple-model ABC sampler: this algorithm falls into the category of *anytime algorithms* [30], because it returns a valid approximate solution (i.e. approximate posterior) even if it is interrupted before the termination criterion is reached.

6 Summary

The paper presents a robust method for localisation of a biochemical source. In the absence of an accurate model of the measurement likelihood, a likelihood-free approximate Bayesian computation (ABC) approach is adopted. Furthermore, since there is no universal dispersion model applicable for all situations (terrain, meteorological conditions), a multiple-model approach is proposed, whereby all candidate models are active in parallel and assigned the probability of being correct. The proposed method computes adaptively

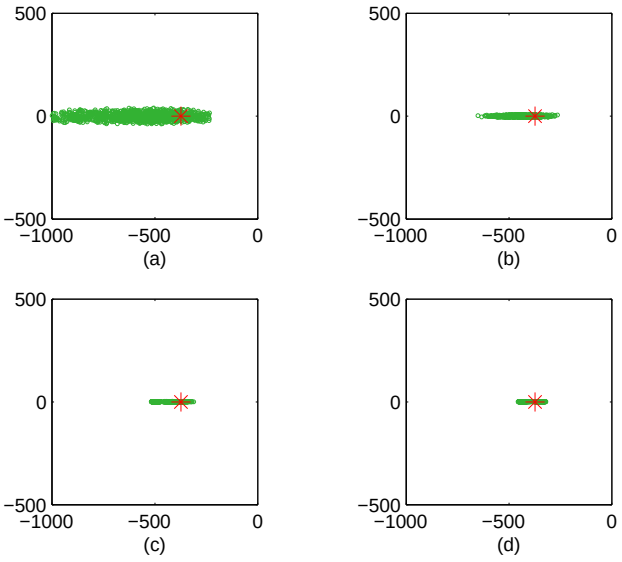

Figure 6: Scatter plots of random samples $\mathcal{L}_m^t$ in the (x, y) plane (dataset 2). The samples approximate the posterior density $p(\ell|\mathbf{z})$. Scatter plots shown after iteration: (a) $t = 3$, (b) $t = 6$, (c) $t = 9$, (d) $t = 12$. The true source location is marked with a red asterisk.

the tolerance levels which are required for ABC sampling. The method has been tested using two experimental datasets (one in the open terrain, the other with obstacles) using three dispersion models (two Gaussian plume models and the stretch exponential model). Source localisation was very accurate using the open terrain dataset. The dataset with obstacles presented a significant challenge, resulting in a bi-modal posterior distribution of source location. Future work will investigate other dispersion models and the performance of source localisation using binary sensors.

References

[1] R. J. Kendall, S. M. Presley, G. P. Austin, P. N. Smith (Eds.), Advances in Biological and Chemical Terrorism Countermeasures, CRC Press, 2008.

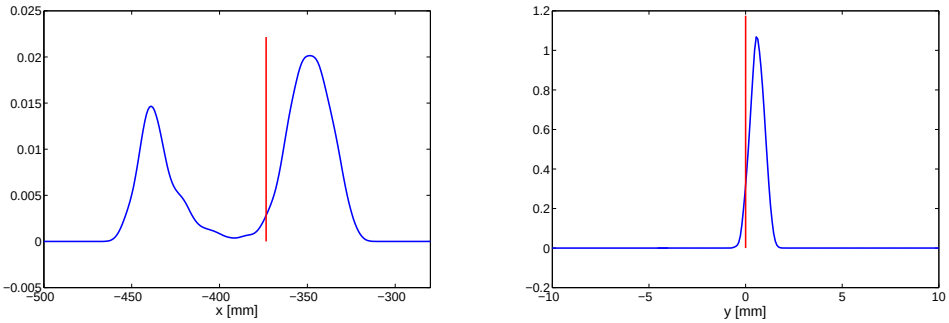

(a) (b)

Figure 7: Estimated posterior density of source location (dataset 2): (a) x axis, (b) y axis. The true values are indicated by the vertical solid lines.

[2] S. P. Arya, Pollution meteorology and dispersion, Oxford University Press, 1998.

[3] N. Holmes, L. Morawska, A review of dispersion modelling and its application to the dispersion of particles, Atmospheric Environment 40 (2006) 5902–5928.

[4] E. Makalic, D. F. Schmidt, A.-K. Seghouane, A tutorial on model selection, in: R. Chellappa, S. Theodoridis (Eds.), Signal processing theory and machine learning, Vol. 1 of Academic press library in signal processing, Elsevier, Oxford, UK, 2014, Ch. 25, pp. 1415–1452.

[5] E. Yee, Inverse dispersion for an unknown number of sources: Model selection and uncertainty analysis, ISRN Applied Mathematics 2012 (2012) 20, article ID 465320.

[6] J. Matthes, L. Gröll, H. B. Keller, Source localization by spatially distributed electronic noses for advection and diffusion, IEEE Trans. Signal Processing 53 (5) (2005) 1711–1719.

[7] L. C. Thomson, B. Hirst, G. Gibson, S. Gillespie, P. Jonathan, K. D. Skeldon, M. J. Padgett, An improved algorithm for locating a gas source using inverse methods, Atmospheric environment 41 (2007) 1128–1134.

[8] A. Keats, E. Yee, F.-S. Lien, Bayesian inference for source determination with applications to a complex urban environment, Atmospheric Environment 41 (3) (2007) 465 – 479.

[9] R. Humphries, C. Jenkins, R. Leuning, S. Zegelin, D. Griffith, Atmospheric tomography: a Bayesian inversion technique for determining the rate and location of fugitive emissions, Environmental Science and Technology 46 (3) (2012) 1739–1746.

[10] M. Ortner, A. Nehorai, A. Jeremic, Biochemical transport modeling and Bayesian source estimation in realistic environments, IEEE Trans. Signal Processing 55 (6) (2007) 2520–2532.

[11] I. Senocak, N. W. Hengartner, M. B. Short, W. B. Daniel, Stochastic event reconstruction of atmospheric contaminant dispersion using Bayesian inference, Atmospheric Environment 42 (33) (2008) 7718 – 7727.

[12] J. A. Hoeting, D. Madigan, A. E. Raftery, C. T. Volinsky, Bayesian model averaging: A tutorial, Statistical science 14 (4) (1999) 382–417.

[13] J.-M. Marin, P. Pudlo, C. P. Robert, R. J. Ryder, Approximate Bayesian computational methods, Statistics and Computing 22 (6) (2012) 1167–1180.

[14] J. K. Pritchard, M. T. Seielstad, A. Perez-Lezaun, M. W. Feldman, Population growth of human Y chromosomes: A study of y chromosome microsatellites, Molecular Biology and Evolution 16 (12) (1999) 1791–1798.

[15] R. Macdonald, Theory and objectives of air dispersion modeling, Tech. rep., University of Waterloo (2003).

[16] C. H. Huang, On solutions of the diffusiondeposition equation for point sources in turbulent shear flow, Journal of Applied Meteorology 38 (1999) 250–254.

[17] M. S. P. Kumar, An analytical model for dispersion of pollutants from a continuous source in the atmospheric boundary layer, Proc. R. Soc. A 466 (2010) 383406.

[18] A. Skvortsov, E. Yee, Scaling laws of peripheral mixing of passive scalar in a wall-shear layer, Phys. Rev. E 83 (3).

[19] G. T. Csanady, Turbulent Diffusion in the Environment, D. Reidel Publishing Company, Boston, USA, 1973.

[20] J. M. Stockie, The mathematics of atmospheric dispersion modeling, SIAM Review 53 (2011) 349–372.

[21] P. Marjoram, J. Molitor, V. Plagnol, S. Tavare, Markov chain Monte Carlo without likelihoods, PNAS 100 (26) (2003) 15324–15328.

[22] S. A. Sisson, Y. Fan, M. M. Tanaka, Sequential Monte Carlo without likelihoods, PNAS 104 (6) (2007) 1760–1765.

[23] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, M. P. H. Strumpf, Approximate Bayesian computation scheme for parameter inference and model selection in dynamic systems, J. R. Soc. Interface 6 (2009) 187–202.

[24] P. D. Moral, A. Doucet, A. Jasra, An adaptive sequential Monte Carlo method for approximate Bayesian computation, Statistics and Computing 22 (5) (2012) 1009–1020.

[25] O. Cappe, A. Guillin, J. M. Marin, P. C. Robert, Population Monte Carlo, Journal of Computational and Graphical Statistics 13 (4) (2004) 907–929.

[26] C. Musso, N. Oudjane, F. LeGland, Improving regularised particle filters, in: A. Doucet, N. deFreitas, N. J. Gordon (Eds.), Sequential Monte Carlo methods in Practice, Springer-Verlag, New York, 2001, Ch. 12.

[27] E. Yee, R. M. Gailis, A. Hill, T. Hilderman, D. Kiel, Comparison of wind-tunnel and water-channel simulations of plume dispersion through a large array of obstacles with a scaled field experiment, Boundary-Layer Meteorology 121 (2006) 389–432.

[28] B. W. Silverman, Density estimation for statistical and data analysis, Chapman and Hall, 1986.

[29] B. Ristic, A. Gunatilaka, R. Gailis, Achievable accuracy in parameter estimation of a Gaussian plume dispersion model, in: Proc. IEEE Workshop Statist. Signal Proc., Gold Coast, Australia, 2014.

[30] S. Zilberstein, Using anytime algorithms in intelligent systems, AI Magazine 17 (3) (1996) 73–83.